

Real-Time Calibration of the Murchison Widefield Array

D. A. Mitchell, L. J. Greenhill, R. B. Wayth, R. J. Sault, C. J. Lonsdale,
R. J. Cappallo, M. F. Morales, and S. M. Ord

Abstract—The interferometric technique known as peeling addresses many of the challenges faced when observing with low-frequency radio arrays, and is a promising tool for the associated calibration systems. We investigate a real-time peeling implementation for next-generation radio interferometers such as the Murchison Widefield Array (MWA). The MWA is being built in Australia and will observe the radio sky between 80 and 300 MHz. The data rate produced by the correlator is just over 19 GB/s (a few Peta-Bytes/day). It is impractical to store data generated at this rate, and software is currently being developed to calibrate and form images in real time. The software will run on-site on a high-throughput real-time computing cluster at several tera-flops, and a complete cycle of calibration and imaging will be completed every 8 seconds. Various properties of the implementation are investigated using simulated data. The algorithm is seen to work in the presence of strong galactic emission and with various ionospheric conditions. It is also shown to scale well as the number of antennas increases, which is essential for many upcoming instruments. Lessons from MWA pipeline development and processing of simulated data may be applied to future low-frequency fixed dipole arrays.

I. INTRODUCTION

The Murchison Widefield Array (MWA) is an 80-300 MHz synthesis array that is being built in Western Australia, with construction to be completed in 2010. The shire of Murchison has a quiet radio environment, making it an excellent site for this and other radio facilities [1]. Each of the 512 antennas will be a 4×4 tile of dipoles. An analogue beamformer at each antenna combines the signals from the 16 dipoles, producing an electronically steerable primary beam with a width of approximately 25° at 150 MHz. When the signals from all antennas are combined, the array will have a synthesized beam with a width of approximately $4.5'$ at 150 MHz. The main science goals of the MWA are the detection of redshifted 21cm emission from the Epoch of Reionization (EoR) [2], transient detection (for example [3]), and remote heliospheric sensing [4]. A schematic of two MWA antenna tiles is shown in Fig. 1.

To make a map of the sky using radio interferometry, one typically builds up an estimate of the 2D Fourier transform of the sky, then applies a Fourier transform to obtain the image. This is known as synthesis imaging, and a good overview of the subject is given in [5]. The measured data that are used to build up the Fourier interference pattern are spatial cross-correlations – or visibilities – that are obtained by correlating voltage streams collected by many pairs of antennas. The visibilities lie in the uvw coordinate frame, where w is the component of the antenna separation vector (or

baseline vector) in the direction of the field center (in units of wavelengths), and u and v are orthogonal coordinates in the plane normal to w (aligned with the corresponding image coordinate axes, l and m). This situation is not naturally a 2D Fourier transform. For small images, w is multiplied by a term that is approximately zero and the 2D nature holds. For large fields this is not the case, but the problem can still be reduced to 2D transforms (a good overview is given in [6]).¹ Post processing typically involves calibrating the visibilities, gridding them onto the uv -plane to form a regularly sampled interference pattern, and then applying a 2D FFT to form an image. Techniques such as self-calibration can then be used in an attempt to improve the calibration by iterating back and forth between the visibilities and the image.

For a number of reasons, MWA visibilities cannot be processed in this way. Many of these effects are common to all low frequency arrays, and are described in detail in [7] and [5]. Each antenna has a different direction-dependent response over the field of view, which cannot simply be divided out. These response patterns may also change significantly over the course of an observation. Furthermore, the ionosphere causes direction-dependent phase shifts that effectively change the position and polarization state of sources during an observation. These effects mean that we cannot make a fully calibrated interference pattern in the standard way. What we can do is use the measured visibilities to iteratively fit ionospheric phase shifts and antenna gains towards many bright catalogue sources, and store these fits to aid deconvolution and resampling processes after the images have been made. These measurements are the focus of this paper, but before they are discussed some of these challenges will be reviewed more closely.

Unlike the radio sky at higher frequencies, which appears sparsely populated at the sensitivity levels of modern instruments, the sky to be observed by the MWA is full of sources. The high density of sources and large angular resolution of the array will result in images that are confusion limited, with significant flux coming from background galaxies in each synthesized beam, as well as from the sidelobes of other sources in the primary beam of each antenna. The sky is also quite complex. There is significant emission on many angular size scales, from compact extragalactic sources and pulsars to diffuse galactic synchrotron radiation [7]. The latter

¹The MWA will produce *snapshot* images, which means each antenna pair contributes a single visibility to each image. The MWA visibilities will be approximately coplanar, so a 2D Fourier relationship will hold even for wide-field images.




Fig. 1. Schematic diagram of an MWA interferometer baseline and sky. Each MWA receiving element is a phased tile comprising a 4×4 grid of crossed dipoles (vertical planar bowtie structures). The main response lobes are steered electronically by beam-formers (BF) to establish the instrument field of view. BF output signals are sampled at baseband and filtered digitally in receiver electronics (RX) and correlated to provide cross-power spectra. In addition to compact objects (e.g., quasars and pulsars), the MWA sky includes foreground emission from the galaxy, which is chiefly synchrotron emission that is linearly polarized and traces turbulent and ordered magnetic fields in the interstellar medium (\vec{B}). Magnetized plasma from solar coronal mass ejections or other activity move outward through the heliosphere. This medium and the Earth ionosphere induce primarily refractive shifts in source positions (higher order distortions in the observing passband are small due to the limited extent of the array) and time-variable Faraday rotation of polarization along different lines of sight. The calibration scheme described here enables polarization-sensitive imaging of the diffuse galactic and compact source populations and assembly of a sky model. Knowledge of the compact background sources is “confusion limited” (i.e., constrained by multiplicity of sources within one instrument resolution element and contamination of each image pixel by sidelobes of distant sources). Confusion increases with image sensitivity and the size of a resolution element.

experiences Faraday rotation and depolarization in the ionized interstellar medium, and as a result exhibits a strong linearly polarized component that is highly dependent on position and frequency [8]. This polarized all-sky signal is at least a few orders of magnitude brighter than the unpolarized EoR signal, and a fully-polarized calibration formalism, such as the Hamaker-Bregman-Sault measurement equation [9], is required to reduce contamination from the spatially structured, linearly polarized emission from the galaxy.

If the raw interferometer data are stored for offline processing, iterative calibration and deconvolution algorithms can be used to address many of the problems described below (see,

for example, [10] and [11]). However, it is impractical to store the 19 GB s^{-1} data stream coming from the MWA correlator,² and the MWA will store images. This means that much of the calibration must take place in real time, before or during the imaging process. At the heart of the calibration system for the MWA is the calibrator measurement loop (CML), which measures apparent angular offsets induced by the ionosphere and the system gain toward known compact astronomical sources across the sky. These measurements are used to fit models of the ionosphere and instrument response, and support subtraction of strong sources that limits sidelobe contamination during calibration. As in [12] and [13], measuring and subtracting the contribution of each source is carried out sequentially, so that the stronger sources are removed before measurements of weaker sources are made. In this paper we do not consider multivariate fits of parameters for all of the calibrator sources simultaneously, as described in [13], but we will discuss it briefly in section IV.

Estimation of calibration parameters is greatly over constrained due to the large number of antenna pairs (1.31×10^5). On the other hand, the wide-field nature of the instrument and real-time computing requirement pose challenges, the most important of which are listed below.

- 1) *Direction-dependent gain and polarization response.* Each MWA receiving element is a 4×4 array of fixed crossed dipoles (Fig. 1). The phased beams are steered with an analogue beamformer, which will typically be updated every 5 to 10 minutes to compensate for rotation of the Earth. The common approach of assuming that the polarized receptors are orthogonal over the field of view with a small amount of direction-dependent leakage cannot be used. The direction-dependent instrumental polarization of the antenna beams will be significant, and it will be measured along with the direction-dependent gain using many calibrator sources spread over the entire sky. These measurements will be repeated as the field of interest moves across the antenna beams.
- 2) *Confusion.* Since the MWA’s primary beams cover such a large section of the sky, each field mapped by the MWA will contain hundreds of relatively bright sources. To calibrate the array, we require accurate flux density measurements of known sources. Such measurements can be corrupted by faint sources within the synthesized beam of the array (“confusion”) as well as the sidelobes of brighter sources outside the synthesized beam (“sidelobe contamination”).

Confusion and sidelobe contamination can arise from both compact and large-scale sources such as extragalactic radio galaxies and galactic synchrotron emission respectively. The Galactic synchrotron, in particular, has a polarized component and structure on many spatial scales. Since the interferometer baselines will respond to large-scale structure differently depending on their

²The $N_a = 512$ antennas lead to $N_a(N_a - 1)/2 = 130816$ different cross-correlation measurements. They are made in a correlator for 4 different polarization products and up to 3072 different frequency channels every 0.5 seconds. Each visibility is represented by 3 real bytes and 3 imaginary bytes, which leads to $\sim 19 \text{ GB/s}$.

length and orientation, calibration of data that includes bright resolved sources such as the Sun or Galactic plane must be performed carefully.

- 3) *Ionosphere*. The 3-dimensional ionosphere can significantly perturb the waves coming from celestial sources. The maximum antenna separation of the MWA is short enough that, for a given source and during normal ionospheric conditions, all of the antennas have approximately the same line of sight through the ionosphere. This assumption will be used throughout. Under these conditions there will be no defocusing and the ionosphere can be described by a two-dimensional phase screen that makes sources appear to move away from their true positions. (Faraday rotation of incident polarization is a second ramification, but this will be considered in a future paper. We limit consideration here to ionospheric calibration using unpolarized sources, which cannot be used to calibrate polarization position angles [9], [14].)
- 4) *Real-time data reduction*. The real-time nature of the MWA means that compute-intensive processes need to be avoided where possible. Unfortunately, this means that many of the promising techniques currently being investigated, such as iterative self-calibration and deconvolution algorithms [10], and some wide-field imaging algorithms [6], cannot currently be implemented in real time. They either cannot be used at all or need to be approximated.

Calibration occurs in a back-end known as the real-time system (RTS), which consists of a visibility integrator (time and frequency), the CML, and an imaging pipeline. These tasks run sequentially, and as mentioned later the processing load is split over frequency. The imaging pipeline incorporates gridding, imaging FFTs, correction for ionospheric and wide-field distortion of the sky, Stokes conversion of images,³ and astronomical coordinate conversion. MWA primary science drivers require that the time and frequency resolution are sufficient to ensure that stationary signals from sources throughout the antenna field of view are coherent for the highest frequencies and longest antenna separations, where interference fringe phases vary most rapidly with time.

Most of the key elements of the real-time calibration system have been coded and are regularly tested with simulated data, as described in section V. To date, the tests have focused on unpolarized cosmic signals, but the response of the instrument polarizes the signals during reception, and the processing employs the fully-polarized description discussed in the ensuing chapters. Apart from tolerance testing – to determine optimal bandwidths, number of calibrators, etc. – and algorithm development for real-time operation, the main piece of outstanding work is the incorporation of polarized calibrators into the system to fully constrain primary beam models and the Faraday rotation component of the ionosphere.

While the discussion and examples given below focus on the

MWA, the techniques are applicable to other low-frequency array projects, such as the SKA Molonglo Prototype (SKAMP), the Long Wavelength Array (LWA), and the Low Frequency Array (LOFAR). That said, the instantaneous synthesized beam of the MWA does make it particularly well suited to this type of processing. We will not go into specific details on how the techniques can be optimized, which at any rate will be different for the different arrays, and refer to papers such as [16] for an overview of the power of large- N radio arrays.

After outlining the assumptions and mathematical model in the next section, the steps in the CML are discussed in more detail in section III, followed by a discussion on algorithm convergence and performance. We then finish with an analysis of peeling simulated MWA data in section V.

II. THE VISIBILITIES

In [14], Hamaker presents a matrix version of self-calibration that leads to a straightforward procedure for estimating each antenna's polarized response to a calibrator source. In this and the next section, the mathematical formalism described in [14] is used to build up the planned implementation for the MWA. As in [14], bold uppercase variables will represent matrices, bold lowercase variables will represent column vectors and a dagger (\dagger) will denote a Hermitian transpose. The input to the CML is a new set of visibilities, measured every Δt seconds and averaged into M frequency channels of width Δf Hz. For the MWA, the data from the correlator are averaged over $\Delta t = 8$ seconds and $\Delta f = 40$ kHz, with $M = 768$, and then sent straight to the CML. The cadence time is set to 8 seconds in order to oversample the time-varying ionosphere, which fluctuates on timescales shorter than a minute.

Suppose that the visibilities can be approximated by the superposition of N_c unresolved calibrator sources that suffer negligible smearing over Δf and Δt ,⁴ and some additive noise from various sources including thermal system noise, confusion from sources such as background radio galaxies, and sidelobes of extended emission and weak point sources. Consider the contribution of one of these calibrator sources, c , to the visibility measured by antennas j and k in the band centered at f Hz. Let the column vector, $\mathbf{r}_{j,c,f}$, contain the response of the two orthogonal polarized components of receiving system j (in instrumental polarization coordinates). It is equal to the product of a 2×2 Jones matrix, $\mathbf{J}_{j,c,f}$, which contains the complex voltage gain of each polarized receptor to each polarized component of the calibrator signal (including all instrumental effects), and the incident 2×1 signal vector (in sky polarization coordinates), $\mathbf{p}_{c,f}$. The incident radiation can be described by a 2×2 covariance matrix, which contains the flux density of the four polarization

³Stokes parameters describe the polarization state of a signal as an unpolarized component, I, two linearly polarized components, Q and U, and a circularly polarized component, V. They are used extensively in radio astronomy, see, for example, [15] and [5].

⁴Following [5], the fractional bandwidth at 140MHz of $\sim 0.03\%$ causes less than a percent of decorrelation on the longest baselines for sources at the edge of the map. The effect in the image is a broadening of those sources by less than 0.01%. The integration time causes a few 10s of percent of decorrelation on the longest baselines, and these will be processed at a faster cadence (2 seconds) to reduce this decorrelation down to the percent level. One should note, however, that the majority of baselines are very short (baseline density goes as the reciprocal of baseline length squared outside a densely-packed core), and most baselines suffer far less decorrelation.

products, $\mathbf{P}_{c,f} = \langle \mathbf{p}_{c,f} \mathbf{p}_{c,f}^\dagger \rangle$, where angle brackets denote the expectation value. In the absence of any ionospheric effects, and if the Jones matrices are constant over the time interval in which the expectation values are estimated, the visibility matrix measured by baseline jk is

$$\begin{aligned} \mathbf{R}_{jk,c,f} &= \left\langle \left(\mathbf{J}_{j,c,f} \mathbf{P}_{c,f} \right) \left(\mathbf{p}_{c,f}^\dagger \mathbf{J}_{k,c,f}^\dagger \right) \right\rangle. \\ &= \mathbf{J}_{j,c,f} \mathbf{P}_{c,f} \mathbf{J}_{k,c,f}^\dagger. \end{aligned} \quad (1)$$

Define \mathbf{s}_c to be the expected position of source c , and $\mathbf{s}'_{c,f} = \mathbf{s}_c + \delta \mathbf{s}_{c,f}$ to be the apparent position of the refracted source, where $\delta \mathbf{s}_{c,f}$ is a small error in the position estimate and a prime will be used to indicate that a variable has been disturbed by the ionosphere. These position vectors can be expressed as phase shifts, $\phi'_{jk,c,f} = \phi_{jk,c,f} + \delta \phi_{jk,c,f}$, where $\phi_{jk,c,f}$ is given by the dot product of the baseline vector and the calibrator position vector: $2\pi \mathbf{u}_{jk,f} \cdot \mathbf{s}_c$. The model for the ionospherically disturbed visibility matrix of baseline jk is the superposition of the contributions from each of the strong sources,

$$\mathbf{V}'_{jk,f} = \mathbf{N}_{jk,f} + \sum_{c=1}^{N_c} \mathbf{R}_{jk,c,f} \exp\{-i\phi'_{jk,c,f}\}, \quad (2)$$

where $\mathbf{N}_{jk,f}$ contains the noise (thermal and confusion) in each polarization product. The negative sign in the exponent follows the convention adopted in [5], which is also the convention used by software packages such as FFTW (a negative exponent is used when transforming from a real plane to an complex plane).

III. THE CALIBRATOR MEASUREMENT LOOP

The general approach chosen for the CML is similar to the peeling approach suggested in [12] and discussed in [13]. We track a few hundred strong point sources, making continual measurements of antenna primary beams and the refractive effect of the ionosphere in each source direction. For this, we need a list of strong radio sources with known positions and fluxes, which we get from the existing catalogues of the southern sky, such as [17] and [18]. However, some bootstrapping will be required for the fluxes, since no catalogue covers the entire MWA frequency range. At least initially, we will only use unresolved point sources as calibrators, since each can be described by a single Fourier component. The Molonglo Reference Catalogue (MRC) contains 7347 sources with flux densities of or above 1 Jy at 408 MHz in the declination range $\delta = -85^\circ$ to $+18.5^\circ$ [19]. Of these, about 90% show no clear evidence of departure from point sources for the MRC beam of $2'.62 \times 2'.86 \text{ sec}(\delta + 35^\circ.5)$.

The calibrator that contributes the most power to the visibilities is selected, which will typically be much stronger than the superposition of the sidelobes from other sources. A model-dependent phase ramp is fit to the visibilities to estimate the ionospheric refraction in the source direction, and then least squares estimates of the direction dependent antenna gains are made. These estimates are used to create models of the contribution to the visibilities, which are subtracted. This process

is then repeated for the rest of the strong calibrators. After this, one will typically make further ionospheric refraction and primary beam measurements without source subtraction.

The natural dimension over which to parallelize the processing is frequency. The imaging pipeline and much of the calibration system can be run independently for different frequency channels, and one can think of the MWA real-time computer as a system of essentially independent threads that each take a subset of the channels (say 10 consecutive 40 kHz channels) to process with the CML (each thread running on its own compute node).

To help isolate ionospheric refraction phase shifts from instrument phase shifts, the threads will be loosely coupled so that the whole 32 MHz band can be used for the ionospheric phase measurements. While the amount of data passed between threads is small,⁵ the sources will need to be synchronized across threads. This is shown schematically in Fig. 2. There is also the potential for pan-frequency averaging during the antenna gain measurements.

Let the number of consecutive frequency channels processed as a group be K , so that each of the M/K nodes hold a $K\Delta f$ Hz sub-band. The flow of the CML in this framework is as follows.

- *Ranking*: Initial antenna primary beam models are used with a catalogue to rank the calibrators by expected received power given the current location of the source and the pointing direction of the antenna. The catalogue will be generated during the commissioning of the array, and the sources will be regularly surveyed to check for time variability. The sequential nature of the peeling and the need for synchronization across the full 32 MHz band mean that the different beam shapes at different frequencies need to be taken into account when ranking sources. The ranking will be updated at a rate given by the motion of sources through the antenna beams, for instance when the antenna beams are updated.
- *Initial source subtraction*: An estimate of the calibrator summation in (2) is made for each visibility and subtracted from it (“pre-peeling”). The sum is over all of the calibrators to be peeled and the idea is to remove as much sidelobe power from the visibility set as possible. These estimates will usually be made from data measured Δt seconds earlier, and provided that the ionosphere has not moved a source more than a synthesized beam in this time, most of the power will be removed. The strong source measurements will be less vulnerable to the subtraction errors of other sources, so they are made first and then peeled properly before the weak source measurements are made.
- *Loop*: Each thread proceeds down the ranked calibrator list, at each step performing the following tasks. These tasks will be elaborated on in sections III-A through III-D.

⁵The summations in (6) need to be generated in each compute node, each providing the information for a different observing frequency. For each source, a central node needs to gather a few 10s of bytes from each compute node and send two fitted coefficients back.

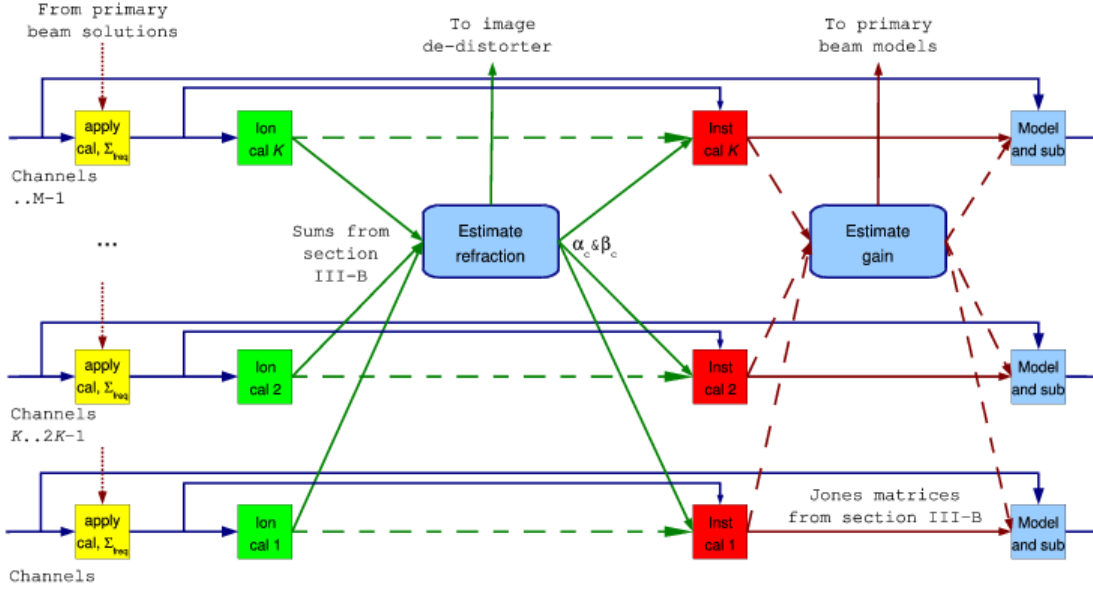


Fig. 2. Parallelization of the CML. The processing for a single calibrator is shown, moving in sequence from left to right. Each horizontal chain is processed by a single compute node, with low-bandwidth intercommunication occurring only at a few key places. Blue lines indicate visibilities, other colors show meta-data, with dashed lines indicating potential data paths.

- 1) *Rotate visibilities and sum over frequency*: Add back into the visibility set the contribution from the current calibrator that was subtracted during the initial source subtraction step. Set the phase center of the visibilities to the estimated calibrator position and average across the K frequency channels, as shown in section III-A. For the MWA, some of the longer baselines will be stored with integration times shorter than Δt . These will also be averaged together at this stage. The averaged visibilities are saved as a new visibility set with reduced thermal noise. Averaging at this point also substantially reduces the number of computing operations performed by the CML.
- 2) *Ionospheric refraction measurements*: Each node creates intermediate sums for the quantities described in section III-B. These sums are gathered from each of the nodes to constrain the pan-frequency ionospheric refraction measurement. For isoplanatic patch sizes of $\sim 4^\circ$ and image sizes of $\sim 30^\circ$, we would need at least 50 or 60 sources in the field of view to describe the phase variation of each patch. Optimally, we would like to oversample these variations by a factor of at least a few. After the fit, the parameters are broadcast back to the nodes and the phase center of each averaged visibility set is shifted to this new position. This is shown by green lines in Fig. 2.
- 3) *Instrumental gain measurements*: Each node then performs the least squares optimization for the polarized voltage gain of each antenna, as described in section III-C. If a higher signal-to-noise ratio is required, the data from each sub-band can be

gathered together for a full-band fit. This is shown by red lines in Fig. 2.

- 4) *Source subtraction*: If the gain and ionosphere measurements pass a set of goodness-of-fit tests, they are used to peel the source from the full resolution visibility set, thus correcting the initial source subtraction that was based on old data. Otherwise the initial source subtraction is repeated. If there are more calibrators in the list, loop back to step 1.

A. Rotate Visibilities and Sum over Frequency

At the start of the loop for calibrator c , the visibilities have had all of the other calibrators peeled off – stronger ones modeled using the current data, weaker ones modeled using less recent data. We will use $\hat{\mathbf{V}}_{jk,c,f_a}$ to designate the peeled visibility matrix for frequency channel f_a . All of the matrices in the sub-band are rotated to be phase centered at the location of the calibrator and summed across frequency,

$$\begin{aligned} \mathbf{V}_{jk,c,f_0}^{(K)} &= \frac{1}{K} \sum_{a=1}^K \frac{\hat{\mathbf{V}}_{jk,c,f_a}'}{\eta_{jk,c,f_a}} \exp\{+i\phi_{jk,c,f_a}\} \\ &\approx \mathbf{R}_{jk,c,f_0} \exp\{-i\delta\phi_{jk,c,f_0}\} + \mathbf{N}_{jk,c,f_0}^{(K)}, \end{aligned} \quad (3)$$

where the superscript (K) indicates an average over K frequency channels, $\eta_{jk,c,f}$ is the product of any undesirable baseline dependent multiplicative factors, such as bandpass shape and decorrelation, and f_0 is the central frequency. As before, the prime indicates that there is an apparent phase shift due to the ionosphere. The error term $\mathbf{N}_{jk,c,f_0}^{(K)}$ contains a rotated version of the noise, $\mathbf{N}_{jk,f}$, and residual sidelobes due to errors in the calibrator source subtraction.

B. Ionospheric Refraction Measurements

The averaged and peeled visibilities given in (3) contain the four instrumental polarization products for calibrator c , with some added noise. These are the data that will be used to fit an ionospheric phase ramp, as described in [20]. Since the visibilities are dominated by flux from a single direction, they can be converted to Stokes parameters, and only the Stokes I visibilities will be considered. Flux densities of calibrator sources are expected to vary smoothly with frequency, and these variations will be catalogued. For the purposes of combining phase data from across the entire observing band, they are divided out so that $|I'_{jk,c,f_0}|$ can be replaced with I_c . The instrument Jones matrices will also vary with frequency, but they are expected to vary in time on scales of minutes, and are assumed to be constant on the ionospheric phase time scales of 10s of seconds. To deal with these variations, equations (1) and (3) show that an estimate, $\mathbf{P}'_{jk,c,f}$, of the sky visibility matrix for source c , can be made using the visibilities calibrated with a recent gain solution

$$\begin{aligned} \mathbf{P}'_{jk,c,f_0} &\approx \mathbf{J}_{j,c,f_0}^{-1} \mathbf{V}'_{jk,c,f_0} \mathbf{J}_{k,c,f_0}^{\dagger-1} \\ &\approx \mathbf{P}_{c,f_0} \exp\{-i\delta\phi_{jk,c,f_0}\} \\ \Rightarrow I'_{jk,c,f_0} &\approx I_c \exp\{-i\delta\phi_{jk,c,f_0}\}. \end{aligned} \quad (4)$$

Since the visibilities are phased towards the source, the sky l and m coordinates for it should be zero, and there should be no phase ramp. Suppose though that, as indicated in (3) and (4), the ionosphere adds a relative phase shift that appears to move the calibrator slightly in the l and m directions, and that the offsets are $\alpha_c \lambda_0^2$ and $\beta_c \lambda_0^2$ respectively, where λ_0 is the wavelength associated with frequency f_0 . Equation (4) becomes

$$\begin{aligned} I'_{jk,c,f_0} &\approx I_c \exp\{-i2\pi(\alpha_c u_{jk,f_0} + \beta_c v_{jk,f_0})\lambda_0^2\} \\ &\approx I_c - i2\pi I_c (\alpha_c u_{jk,f_0} + \beta_c v_{jk,f_0})\lambda_0^2, \end{aligned} \quad (5)$$

where u_{jk,f_0} and v_{jk,f_0} are components of the baseline vector introduced in section I, and the expansion only holds when $(\alpha_c u_{jk,f_0} + \beta_c v_{jk,f_0})\lambda_0^2 \ll 1$. Erickson found the root-mean-square displacement of sources observed with the Clark Lake TPT telescope at 74 MHz to be about 1'.1, [21]. Erickson only considered relatively long-period fluctuations (~ 1 hr), and we anticipate the variations for 8-second periods to be much less than this, with root-mean-square displacements of several arcseconds or less. For arcminute deviations, the expansion in (5) breaks down on the long MWA baselines, so only short baselines are used for the initial fits. We then track the short-period deviations.

Each visibility contains the sum of components from thousands of cosmic sources, with different strengths and random phases, and this sum, along with additive thermal noise from the receiving system, is expected to be very close to normally distributed. This does indeed appear to be the case for the simulated visibilities discussed section V, with and without source peeling. If the noise is independent with variance

σ_{jk,f_0}^2 , and $\Re()$ and $\Im()$ are the real and imaginary operators respectively, the least squares solutions for $I_{jk,c}$, α_c and β_c in (5) are

$$\begin{aligned} I_c &= \frac{\sum_{jkf} \Re(I'_{jk,c,f}) \sigma_{jk,f}^{-2}}{\sum_{jkf} \sigma_{jk,f}^{-2}} \\ \alpha_c &= (a_{vv}A_u - a_{uv}A_v) / (2\pi I_c \Delta) \\ \beta_c &= (a_{uu}A_v - a_{uv}A_u) / (2\pi I_c \Delta), \end{aligned} \quad (6)$$

where

$$\begin{aligned} a_{uv} &= \sum_{jkf} u_{jk} v_{jk,f} \sigma_{jk,f}^{-2} \lambda_0^4 \\ A_u &= -\sum_{jkf} u_{jk} \Im(I'_{jk,c,f}) \sigma_{jk,f}^{-2} \lambda_0^2 \\ \Delta &= a_{uu}a_{vv} - a_{uv}^2 \end{aligned} \quad (7)$$

The sums on the right-hand side of (6) can be partially generated in each compute node, and then α_c and β_c generated after the partial sums have been gathered together.

C. Instrumental Gain Measurements

If a sufficiently reliable ionospheric refraction measurement has been made, then the phase center of the averaged visibility set can be moved to the new position, and (3) can be corrected:

$$\begin{aligned} \mathbf{V}_{jk,c,f_0}^{(K)} &\approx \mathbf{V}'_{jk,c,f_0} \exp\{+i\delta\phi_{jk,c,f_0}\} \\ &\approx \mathbf{R}_{jk,c,f_0} + \mathbf{N}_{jk,c,f_0}^{(K)}. \end{aligned} \quad (8)$$

To estimate the gain of an antenna towards calibrator c , a similar approach to those described in [22] and [23] will be taken, where one uses all of the visibilities that were measured with an antenna to constrain a simple model of its gain. However, a matrix form from [14] will be followed here. This is a matrix least-squares problem in which one searches for the matrices $\mathbf{J}_{j,c,f_0}^{-1}$ that minimize

$$\sum_{j=1}^{N_a} \sum_{k \neq j}^{N_a} \left\| \mathbf{P}_{jk,c,f_0}^{(K)} - \mathbf{J}_{j,c,f_0}^{-1} \hat{\mathbf{V}}_{jk,c,f_0} \mathbf{J}_{k,c,f_0}^{\dagger-1} \right\|_F^2, \quad (9)$$

where N_a is the number of antenna tiles, $\|\mathbf{A}\|_F^2$ is the squared Frobenius norm of a matrix \mathbf{A} , equal to the trace of $\mathbf{A}\mathbf{A}^\dagger$, and $\hat{\mathbf{V}}_{jk,c,f_0}$ is a model of the measured visibility matrix for source c . In early investigations solutions to (9) were unstable when there were significant sidelobes from other sources, and rearranging (9) had more robust solutions:

$$\sum_{j=1}^{N_a} \sum_{k \neq j}^{N_a} \left\| \mathbf{V}_{jk,c,f_0}^{(K)} - \mathbf{J}_{j,c,f_0} \hat{\mathbf{P}}_{c,f_0} \mathbf{J}_{k,c,f_0}^{\dagger} \right\|_F^2. \quad (10)$$

For each antenna (10) has the analytic solution

$$\mathbf{J}_{j,c,f_0} = \left(\sum_{k,k \neq j}^{N_a} \mathbf{V}_{jk,c,f_0}^{(K)} \mathbf{J}_{k,c,f_0} \hat{\mathbf{P}}_{c,f_0}^\dagger \right) \times \left(\sum_{k,k \neq j}^{N_a} \hat{\mathbf{P}}_{c,f_0} \mathbf{J}_{k,c,f_0}^\dagger \mathbf{J}_{k,c,f_0} \hat{\mathbf{P}}_{c,f_0}^\dagger \right)^{-1} \quad (11)$$

Equation (11) can be used to estimate a new set of Jones matrices for each calibrator (or to update matrices that are restricted to moving more slowly). Once the matrices have converged they can be left to track slow changes in the gain for each calibrator.

Finally, it is worth pointing out that often, much of the information in the Jones matrices is known. For example, the transformation from the sky polarization coordinates to instrumental polarization coordinates may be known to high precision. If this is the case, it is straightforward to modify (10) so that the other matrices absorb the known information, leaving only the unknown quantities for the fit.

D. Source Subtraction

The dynamic range of power levels expected from the hundred or so strongest calibrators will be many orders of magnitude. This is the driving factor behind the sequential approach discussed above; stronger sources are subtracted before measurements are made for the weaker sources. A preliminary investigation along the lines of [24] indicates that the deepest we will be able to peel or clean is to a sidelobe noise floor of several hundred mJy. At this point, all remaining point sources will be less than 5 sigma above the noise. This suggests that there should be on the order of a few hundred calibrators available, if we peel as deeply as possible.

The subtraction step is straightforward, since all of the CML sources have had their apparent positions measured, as described in section III-B, and have been calibrated using the antenna gain models given in (11). Before peeling, however, the models need to be multiplied by the complex $\eta_{jk,c,f}$ factors. These were divided out when averaging to a central frequency in section III-A.

It may be that subtracting 100% of each source is not optimal, since noise in each measurement is added back into the visibilities. Some modifications that work more like the CLEAN algorithm (see [25]) are being considered, and will be tested soon.

IV. DISCUSSION

Performance

The relatively short cadence time of the RTS means we need to make compromises when designing real-time peeling process. One of these compromises will be the number of sources peeled every 8 seconds. Several methods for calibrating antenna gains are compared in [26]. These are alternatives to the technique discussed in section III-C. The number of complex multiplications required by most of the techniques scales as N_a^3 , where N_a is the number of antennas (or, more generally, the number of receivers being correlated to form

visibilities). The technique that required the fewest computations was the logarithmic least squares (LOGLS) algorithm, which scaled as N_a^2 . The numbers given in [26] for a single polarization version are $2N_a^2$ multiplications with an additional $16N_a^2$ for weighting.

How does this compare to the algorithm discussed in section III-C? In our application we measure two polarizations, so if we consider each antenna as 2 polarized receptors the number of multiplications for the LOGLS algorithm is $2(2N_a)^2$. There will probably be another factor of 2 since correlations between the receptors on the same antenna will most likely need to be considered.⁶ The number of complex multiplications used to determine the calibration solutions in section III-C is $O(24N_a^2)$, with an additional $O(12N_a^2)$ for weighting. This efficiency appears to be about as high as one might reasonably expect to achieve.

We are not just dealing with a single source, however. The CML needs to pre-peel all of the calibrator sources, and then, for each source, unpeel, rotate all of the visibilities, solve for the ionospheric offset, solve for the antenna gains, and peel. Table I shows the approximate number of floating-point operations used in each of these steps. These numbers were obtained by listing the main operations in the inner loops of the routines and multiplying each by our best estimate of the associated floating-point operations. These numbers are not exact; they are provided as a rough indication of where processing time will be spent. Also, the number of sources processed by each of the routines need not be the same. For example, we might peel and make gain measurements for 50 sources, but make ionospheric measurements on a few hundred more (which requires only the third and fourth rows).

TABLE I
APPROXIMATE FLOATING-POINT OPERATIONS REQUIRED FOR EACH SOURCE IN THE CML (WITH 512 ANTENNA TILES AND A SINGLE FREQUENCY CHANNEL).

Routine	Floating-point operations (millions)
peeling (applied 3 times)	31
rotate and accumulate	26
ionospheric sums and rerotation	21.5
measure tile gains	37
total	$O(180)$

While the algorithm scales well, there is still on the order of 180 million floating-point operations required during each 8 second calibration cycle. For the whole array, some of the rows in the table need to be multiplied by the number of frequency channels (768), while others need to be multiplied by the number of frequency sub-bands (~ 50). They also need to be multiplied by the number of sources, as discussed in the previous paragraph. We anticipate a few trillion floating-point operations for the CML over the 8 seconds. This does not include various overheads such as memory access that will increase the number of operations by a factor of a few.

⁶The authors of [26] note that LOGLS is not easily generalized to a dual-polarized telescope array, but, for the sake of comparison, suppose that such a generalization exists.

Options for Reducing the Load

Each ionospheric refraction measurement can use data from every baseline, polarization and frequency. That is over 1.5 gigasamples for a single 2D offset. For sources that are only used for ionospheric measurements, all of the processing listed in table I will be reduced if we only use a subset of the visibilities. Most of the baselines are short, and many of those will be redundant. Furthermore, long baselines do not see as much of the extended galactic structure as short baselines, and they measure the apparent offset with higher angular resolution. Initial investigations suggest that we may be able to ignore more than 99% of the short baselines for the stronger sources.

Since the tile gains are changing slowly, we do not necessarily need to make measurements for every source every 8 seconds. This can be exploited by making measurements for the strongest sources every 8 seconds, and cycling through subsets of the other sources. This way we still have gain measurements distributed across the sky every few minutes when we make fits for the tile beams, but we only run the full peel algorithm on 20 or 30 sources at a time.

Rather than solving for direction-dependent parameters sequentially, one could fit for all of the calibrator sources simultaneously. In fact, the peeling algorithm is really just a robust and efficient method for reducing the number of unknowns and finding the multivariate solutions. One can also reduce the degrees of freedom by changing the fit parameters to quantities that do not change (or change slowly) with time or frequency, and solve using multiple snapshots, as discussed in detail in [13]. For the MWA, we hope to be able to use slowly varying direction-dependent dipole gains and phases to describe the primary beam of each tile. Once we have high quality measurements of our tile beams in the field, we will look more closely at fitting these dipole parameters directly. However, source subtraction will then be based on these dipole fits, not on direct tile gain measurements.

Foreground Subtraction

One of the primary challenges in the search for a signature from the EoR is that of foreground subtraction. At best the signal will be several orders of magnitude weaker than the galactic foreground and it is important that we understand the nature of the residuals from the calibration and peeling process. One of the drawbacks of peeling strong sources in real time and then averaging the resulting images together is that any residuals are also averaged into the mix. All of the calibration data will be stored in a database for use in offline processing, and peeling errors can be assessed and reduced at that stage. However, if there is any concern that residuals from the peeling process might mimic the EoR signal then peeling can be used for calibration only, and images formed from unpeeled visibilities.

There is significant effort going into techniques for removing foregrounds during offline processing. These include techniques that exploit spectral differences in the foregrounds and the EoR signature (see for example [27] and the discussion and references in [28]), and the rotation measure synthesis

techniques discussed in [29] & [30]. Direction-dependent deconvolution techniques such as the one described in [10] will also be essential for imaging.

Radio Frequency Interference

Of concern for any telescope operating at MWA frequencies is radio frequency interference (RFI). Even though the Murchison site is extremely radio quiet [1], the array will still have to deal with some RFI. This includes communication and military satellite signals, reflections of FM radio broadcasts, and natural interference such as lightning. Due to the low spectral occupancy of the RFI, the high quality polyphase filter banks used to isolate frequencies, and the campaign-mode operation of the array, we will adopt the traditional strategy of flagging and ignoring contaminated data before imaging. Missing frequency or time samples can be accounted for in the weighting of the various least-squares algorithms.

V. EXAMPLES

To test the CML, we generate simulated visibilities using MAPS, the MIT Array Performance Simulator [31]. Briefly, MAPS models all physical processes of a radio interferometer from the ionosphere, through the analog beamforming in the array tiles to time and frequency averaging in a correlator. MAPS uses polarized receptors and a polarized model sky to generate model visibilities in linear or circular polarization products. The model sky includes a large-scale diffuse component, based on [32] with additional polarized flux, plus a catalog of southern point sources based on [17] and [18].

The tile beams for MWA simulations consist of 16 dual polarization receptors in a 4x4 array. Receptors can have non-equal complex gains, which allows us to include realistic differences in the tile primary beams that we might expect due to the analog parts of the system. In the examples that follow, complex Gaussian noise was added to the gain of each receptor.

MWA's synthesized beam has modest resolution, which allows us to use a very realistic full-sky model as an input to MAPS by simultaneously including large-scale diffuse structure and thousands of point sources. As such, the simulations show a large range of correlated power depending on baseline length. MAPS also implements a model ionosphere to change the relative path length for each baseline in each look direction. The model uses the International Reference Ionosphere [33] for large scale structure, and a Kolmogorov spectrum to add turbulence at smaller scales. The turbulent power was a set fraction of the total power – about 2% – and was repeated over the sky in patches of about 1 square degree. Phase variations due to traveling ionospheric disturbances are not included in these simulations, but they have recently been added to MAPS. Finally, 200K thermal noise was added to the visibilities. In this paper, while the description is fully polarized and there is significant instrumental polarization, we only use unpolarized input since testing of the MAPS software with polarized input is incomplete. Polarized sources will be needed to fully describe the polarized response of each antenna and the Faraday rotation state of the ionosphere [9], [14].




Fig. 3. Uncalibrated Stokes I images made after peeling. Shown are $45^\circ \times 45^\circ$ images, however sources are peeled from anywhere in the sky, depending on their apparent strength. As more sources are peeled, weaker point sources are revealed until, in this field, the galactic center and its sidelobes dominate. Gray-scales are set by the minimum and maximum pixel values, which have units of Jy/beam.

MAPS was used to generate a series of visibility sets at a local sidereal time at the MWA site of 16.5 hours, with the antenna beams pointed at the zenith. Fig. 3 shows three images created after the CML had converged. The only difference in the processing of each image is the number of sources that were peeled. Fig. 3a shows the case for no peeling, where a few strong point sources dominate the image. There is also a slight hint of diffuse galactic structure in the lower left. As background radio sources are peeled, weaker sources and the diffuse galactic foreground become apparent. In Fig. 3c all of the stronger sources have been analyzed and peeled, and the galactic center and its wavy sidelobes completely dominate.

Fig. 4a shows the image noise RMS as a function of the number of sources being peeled (in one of the instrument polarizations). The solid curve is from the simulations described above, the points are for images generated using the point sources only (no galactic emission or ionosphere). After the first few strong sources are peeled away, the image RMS becomes dominated by the galaxy. To reduce the image RMS beyond this point, more sophisticated foreground subtraction algorithms need to be employed, such as those discussed in section IV (one should note that EoR observations will not be made in a field that contains the galactic center). Fig. 4b shows the RMS of the tile gain error towards the strongest calibrator as a function of the number of sources being peeled. Here the curve and the points converge at the same rate, indicating that for this source the algorithm is not limited by the ionosphere or the galactic emission. In these examples, we have weighted down visibilities from short baselines, since they see all of the galactic emission (Fourier components of angular features larger than the reciprocal of the baseline length cancel destructively, so we give more weight to the longer baselines).

Also of importance is how quickly convergence is achieved. At the end of each iteration (for the MWA, each successive iteration is associated with a new set of visibilities, i.e., they occur 8 seconds apart), the new solutions are used to update the old solutions. Since the data are noisy, a weighted average




Fig. 4. Convergence as a function of the number of source peeled: a) RMS of the image noise (Jy/beam), b) RMS of the gain error for the strongest source. The solid curves show full simulations that include point sources, galactic emission and ionospheric effects, while the points represent simulations that only had point sources. The similarity of the gain error curves suggests that the algorithm performs well in the face of an ionosphere and galactic emission.

of new and old solutions are used to set the weights for the next iteration. In these simulations, the weights used for each source were based on the estimated contribution of the source to the visibilities. In other words, strong sources are updated quickly (the strongest being allowed to update by 50% with each iteration), while weak sources are adapted more slowly, some by only a few percent each time. Fig. 5 shows the mean gain errors for the 5 strongest sources, normalized by the gain towards each source. The errors are reduced until convergence is limited by a local minimum in the minimization process. The relative errors appear to converge to the same level, which has not been investigated.

The properties and limits of the convergence shown in Fig. 5 is of the utmost importance. In the current simulations, typically a few tens of calibrators converge well, but some weaker sources can be significantly affected by sidelobes and




Fig. 5. Convergence as a function of time. Shown is the mean relative gain error for the 5 strongest sources. The solid curves represents full simulations that include point sources, galactic emission and ionospheric effects, and the points show simulations that only had point sources. Again, the convergence is not being limited by the ionosphere or the galactic emission.

converge to local minima of (10). A quantitative analysis of the situation is currently underway, and several options are being tested. These include investigations of the calibration bandwidth (sidelobes of distant sources will decorrelate more as the bandwidth of the sub-bands is increased), large pre-peels (approximate subtraction of extra catalogue sources helps reduce the size of local minima), and incorporating all-sky ionospheric and primary beam models to improve peels for weak sources. The primary beam fits are seen to work better when the calibrators are not just ranked by their apparent flux density, but also their position in the beams, so that a larger fraction of them lie in the structured antenna sidelobes.

VI. SUMMARY

We have described a general approach for making measurements of strong point sources that can be used in the calibration process of wide-field, low-frequency radio arrays. This approach has been adopted for the MWA, and there is an ongoing effort to develop the required software, as well as to understand the benefits and limitations of the approach.

We have used simulated visibility data to show that the peeling algorithm works well in situations that are of major concern for future radio telescopes: crowded fields, strong galactic emission, and ionospheric refraction. The algorithm exhibits fast convergence, which is important since sources will be moving in and out of antenna sidelobes and the algorithm needs to be able to keep up with the antenna gain and phase changes, as well as changes in the ionosphere. Critical parts of the process are shown to be computationally efficient, and parts of the system lend themselves to significant levels of optimization.

The next step is a detailed analysis of the convergence properties of the algorithms, and a series of tolerance tests to investigate how the algorithms will behave in the various conditions we expect to encounter. This includes observations of weak emission that is masked by significant polarized

diffuse foregrounds, high dynamic range observations close to the sun, and observations in the presence of severe ionospheric conditions, such as during activation and recombination of the ionosphere. As data from the initial deployment of antennas become available in late 2008 and early 2009, we will get a clearer picture of how harmful phenomena such as source variability (due to ionospheric scintillation, for instance) and dipole mutual coupling (which will affect our tile beam models) can be, and these can be worked into the tolerance tests.

ACKNOWLEDGMENTS

We would like to acknowledge the international MWA collaboration for its continued input and support of this work.

REFERENCES

- [1] J. D. Bowman, D. G. Barnes, F. H. Briggs, B. E. Corey, M. J. Lynch, N. D. R. Bhat, R. J. Cappallo, S. S. Doeleman, B. J. Fanous, D. Herne, J. N. Hewitt, C. Johnson, J. C. Kasper, J. Kocz, E. Kratzenberg, C. J. Lonsdale, M. F. Morales, D. Oberoi, J. E. Salah, B. Stansby, J. Stevens, G. Torr, R. Wayth, R. L. Webster, and S. B. Wyithe, "Field Deployment of Prototype Antenna Tiles for the Mileura Widefield Array Low Frequency Demonstrator," *Astron. J.*, vol. 133, no. 4, pp. 1505–1518, Apr., 2007.
- [2] J. D. Bowman, M. F. Morales, and J. N. Hewitt, "The Sensitivity of First-Generation Epoch of Reionization Observatories and Their Potential for Differentiating Theoretical Power Spectra," *Astrophys. J.*, vol. 638, no. 1, pp. 20–26, Feb. 2006.
- [3] N. D. R. Bhat, R. B. Wayth, H. S. Knight, J. D. Bowman, D. Oberoi, D. G. Barnes, F. H. Briggs, R. J. Cappallo, D. Herne, J. Kocz, C. J. Lonsdale, M. J. Lynch, B. Stansby, J. Stevens, G. Toor, R. L. Webster, and S. B. Wyithe, "Detection of Crab Giant Pulses Using the Mileura Widefield Array Low Frequency Demonstrator Field Prototype System," *Astrophys. J.*, vol. 665, no. 1, pp. 618–627, Aug. 2007.
- [4] J. E. Salah, C. J. Lonsdale, D. Oberoi, R. J. Cappallo, and J. C. Kasper, "Space weather capabilities of low frequency radio arrays," in *Proc. SPIE: Solar Physics and Space Weather Instrumentation*, San Diego, 2005, vol. 5901, pp. 124–134.
- [5] A. R. Thompson, J. M. Moran, and G. W. Swenson, Jr., *Interferometry and Synthesis in Radio Astronomy*, 2nd Ed., New York: Wiley, 2001.
- [6] T. J. Cornwell, K. Golap, and S. Bhatnagar, "W Projection: A New Algorithm for Wide Field Imaging with Radio Synthesis Arrays", in *ASP Conf. Series: Astronomical Data Analysis Software and Systems XIV*, San Francisco: Astronomical Society of the Pacific, 2005, vol. 347, pp. 86–90.
- [7] W. C. Erickson, "Long Wavelength Interferometry," in *ASP Conf. Series*, 1999, vol. 180, pp. 601–612.
- [8] M. H. Wieringa, A. G. de Bruyn, D. Jansen, W. N. Brouw, and P. Katgert, "Small scale polarization structure in the diffuse galactic emission at 325 MHz," *Astron. and Astrophys.*, vol. 268, pp. 215–229, 1993.
- [9] R. J. Sault, J. P. Hamaker, and J. D. Bregman, "Understanding radio polarimetry. II. Instrumental calibration of an interferometer array," *Astron. Astrophys. Suppl. Ser.*, vol. 117, pp. 149–159, May, 1996.
- [10] S. Bhatnagar, T. J. Cornwell, K. Golap and J. M. Uson, "Correcting direction-dependent gains in the deconvolution of radio interferometric images," *Astron. Astrophys.*, in press.
- [11] R. J. Nijboer and J. E. Noordam, "LOFAR Calibration," in *ASP Conf. Series: Astronomical Data Analysis Software and Systems XVI*, San Francisco: Astronomical Society of the Pacific, 2007, vol. 376, pp. 237–240.
- [12] J. E. Noordam, "LOFAR Calibration Challenges," in *Proc. SPIE: Groundbased Telescopes*, Glasgow, 2004, vol. 5489, pp. 817–825.
- [13] S. van der Tol, B. D. Jeffs, and A. J. van der Veen, "Self-Calibration for the LOFAR Radio Astronomical Array," *IEEE Trans. Signal Process.*, vol. 55, no. 9, pp. 4497–4510, Sep. 2007.
- [14] J. P. Hamaker, "Understanding radio polarimetry. IV. The full-coherency analogue of scalar self-calibration: Self-alignment,
- [15] M. Born and E. Wolf, *Principles of Optics*, 7th Ed., Cambridge, UK: Cambridge University Press, 1999.

- [16] C. J. Lonsdale, S. S. Doeleman, and D. Oberoi, “Efficient Imaging Strategies For Next-Generation Radio Arrays,” *Experimental Astron.*, vol. 17, no. 1–3, pp. 345–362, June, 2004. dynamic range and polarimetric fidelity,” *Astron. Astrophys. Suppl. Ser.*, vol. 143, pp. 515–543, 2000.
- [17] A. Wright and R. Otrupcek, “Parkes Catalog, 1990, Australia telescope national facility,” *PKS Catalog (1990)*. [Online]. Available: <http://vizier.cfa.harvard.edu/viz-bin/VizieR?-source=VIII/15>
- [18] T. Mauch, T. Murphy, H. J. Buttery, J. Curran, R. W. Hunstead, B. Piestrzynski, J. G. Robertson, and E. M. Sadler, “SUMSS: a wide-field radio imaging survey of the southern sky - II. The source catalogue,” *Mon. Not. Roy. Astron. Soc.*, vol. 342, no. 4, pp. 1117–1130, July, 2003.
- [19] M. I. Large, L. E. Cram, and A. M. Burgess, “A machine-readable release of the Molonglo Reference Catalogue of Radio Sources,” *The Observatory*, vol. 111, pp. 72–75, 1991
- [20] R. J. Sault, “Visibility-Based Position Finder,” *MWA-LFD memo: in prep.*, 2007.
- [21] W. C. Erickson, “Ionospheric refraction in radio source observations at long radio wavelengths,” *J. Astrophys. Astr.*, vol. 5, pp. 55–65, Mar. 1984.
- [22] S. Bhatnagar and R. Nityananda, “Solving for closure errors due to polarization leakage in radio interferometry of unpolarized sources,” *Astron. Astrophys.*, vol. 375, pp. 344–350, 2007.
- [23] D. H. Rogstad. (2005, August). *The SUMPLE Algorithm for Aligning Arrays of Receiving Radio Antennas: Coherence Achieved with Less Hardware and Lower Combining Loss*. Jet Propulsion Laboratory: CA. [Online]. Available: http://ipnpr.jpl.nasa.gov/progress_report/42-162/162B.pdf
- [24] J. J. Condon, “Confusion and Flux-Density Error Distributions,” *Astrophys. J.*, vol. 188, pp. 279–286, Mar. 1974.
- [25] J. A. Högbom, “Aperture Synthesis with a Non-Regular Distribution of Interferometer Baselines,” *Astron. Astrophys. Suppl.*, vol. 15, pp. 417–426, 1974.
- [26] A. J. Boonstra and A. J. van der Veen, “Gain calibration methods for radio telescope arrays,” *IEEE Trans. Signal Process.*, vol. 51, no. 1, pp. 25–38, Jan. 2003.
- [27] M. F. Morales, J. D. Bowman, and J. N. Hewitt, “Improving Foreground Subtraction in Statistical Observations of 21 cm Emission from the Epoch of Reionization,” *Astrophys. J.*, vol. 648, no. 2, pp. 767–773, Sep. 2006
- [28] S. R. Furlanetto, S. P. Oh, and F. H. Briggs, “Cosmology at low frequencies: The 21 cm transition and the high-redshift Universe” *Physics Reports*, 433, no. 4–6, pp. 181–301, Oct. 2006,
- [29] B. J. Burn, “On the Depolarization of Discrete Radio Sources by Faraday Dispersion,” *Mon. Not. Roy. Astron. Soc.*, vol. 133, pp. 67–83, 1966.
- [30] M. A. Brentjens and A.G. de Bruyn, “Faraday rotation measure synthesis,” *Astron. Astrophys.*, vol. 441, pp. 1217–1228, 2005.
- [31] R. B. Wayth, N. D. R. Bhat, R. J. Cappallo, S. S. Doeleman, L. J. Greenhill, J. C. Kasper, C. J. Lonsdale, D. A. Mitchell, and Ord, S. M., “The MIT Array Performance Simulator,” in prep, 2008.
- [32] A. de Oliveira-Costa, M. Tegmark, B. M. Gaensler, J. Jonas, T. L. Landecker, and P. Reich, “A model of diffuse Galactic Radio Emission from 10 MHz to 100 GHz,” *Mon. Not. Roy. Astron. Soc.*, in press, 2008.
- [33] K. Rawer, D. Bilitza, and S. Ramakrishnan, “Goals and Status of the International Reference Ionosphere,” *Rev. Geophys.*, vol. 16, pp. 177–181, 1978.